

#4

Patent
Attorney's Docket No. 030650-074

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of)
Stefan DOBLER et al.)
Application No.: 09/985,844) Group Art Unit: 2641
Filed: November 6, 2001) Examiner: Unassigned
For: METHOD AND DEVICE FOR)
GENERATING AN ADAPTED)
REFERENCE FOR AUTOMATIC)
SPEECH RECOGNITION)



CLAIM FOR CONVENTION PRIORITY

Assistant Commissioner for Patents
Washington, D.C. 20231

Sir:

The benefit of the filing date of the following prior foreign application in the following foreign country is hereby requested, and the right of priority provided in 35 U.S.C. § 119 is hereby claimed:

European Patent Application No. 00123488.9

Filed: November 7, 2000

In support of this claim, enclosed is a certified copy of said prior foreign application. Said prior foreign application was referred to in the oath or declaration. Acknowledgment of receipt of the certified copy is requested.

Respectfully submitted,

BURNS, DOANE, SWECKER & MATHIS, L.L.P.

Date: January 11, 2002

By: Kenneth B. Leffler
Kenneth B. Leffler
Registration No. 36,075

P.O. Box 1404
Alexandria, Virginia 22313-1404
(703) 836-6620

THIS PAGE BLANK (USPTO)





**Europäisches
Patentamt**

**European
Patent Office**

**Office européen
des brevets**



Bescheinigung

Certificate

Attestation

Die angehefteten Unterlagen stimmen mit der ursprünglich eingereichten Fassung der auf dem nächsten Blatt bezeichneten europäischen Patentanmeldung überein.

The attached documents are exact copies of the European patent application described on the following page, as originally filed.

Les documents fixés à cette attestation sont conformes à la version initialement déposée de la demande de brevet européen spécifiée à la page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

00123488.9

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

I.L.C. HATTEN-HECKMAN

DEN HAAG, DEN
THE HAGUE, 18/10/01
LA HAYE, LE



THIS PAGE BLANK (USPTO)



Europäisches
Patentamt

European
Patent Office

Office eur péen
des brevets

#4

Blatt 2 der Besch inigung
Sheet 2 of the certificate
Page 2 de l'attestation

Anmeldung Nr.:
Application no.:
Demande n°: 00123488.9

Anmeldetag:
Date of filing: 07/11/00
Date de dépôt:

Anmelder:
Applicant(s):
Demandeur(s):
Telefonaktiebolaget L M Ericsson (Publ)
126 25 Stockholm
SWEDEN

Bezeichnung der Erfindung:
Title of the invention:
Titre de l'invention:

Reference templates adaptation for speech recognition

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s) revendiquée(s)

Staat:
State:
Pays:

Tag:
Date:
Date:

Aktenzeichen:
File no.
Numéro de dépôt:

Internationale Patentklassifikation:
International Patent classification:
Classification internationale des brevets:

G10L15/06

Am Anmeldetag benannte Vertragsstaaten:
Contracting states designated at date of filing: AT/BE/CH/CY/DE/DK/ES/FI/FR/GB/GR/IE/IT/LI/LU/MC/NL/PT/SE/TR
Etats contractants désignés lors du dépôt:

Bemerkungen:
Remarks:
Remarques:

SEE FOR ORIGINAL TITLE OF THE APPLICATION
PAGE 1 OF THE DESCRIPTION.

THIS PAGE BLANK (USPTO)

EPO - Munich
69
07. Nov. 2000Method and device for generating an adapted reference for
automatic speech recognition

5

BACKGROUND OF THE INVENTION

1. Technical Field

10 The invention relates to the field of automatic speech
recognition and more particularly to generating a reference
adapted to an individual speaker.

2. Discussion of the Prior Art

15

During automatic speech recognition a spoken utterance is
analyzed and compared with one or more already existing
references. If the spoken utterance matches an existing
reference, a corresponding recognition result is output. The
20 recognition result can e. g. be a pointer which identifies the
existing reference which is matched by the spoken utterance.

The references used for automatic speech recognition can be
both speaker dependent and speaker independent. Speaker
25 independent references can e. g. be created by averaging
utterances of a large number of different speakers in a
training process. Speaker dependent references for an
individual speaker, i. e., references which are personalized in
accordance with an individual speaker's speaking habit, can be
30 obtained by means of an individual training process. In order
to keep the effort for the training of speaker dependent
references low, it is preferable to use a single word spoken in
isolation for each speaker dependent reference to be trained.
The fact that the training utterances are spoken in isolation
35 leads to problems for connected word recognition because
fluently spoken utterances differ from utterances spoken in
isolation due to coarticulation effects. These coarticulation
effects deteriorate the accuracy of automatic speech

recognition if speaker dependent references which were trained in isolation are used for recognition of connected words. Moreover, even if connected words have been trained, a user's voice may change, e. g. due to different health conditions, which also deteriorates the accuracy of automatic speech recognition which is based on speaker dependent references. The accuracy of automatic speech recognition is generally even lower if speaker independent references are used, especially when the utterances are spoken in a heavy dialect or with a foreign accent. The accuracy of automatic speech recognition is also influenced by the speaker's acoustic environment, e.g. the presence of background noise or the use of a so-called hands free set.

In order to improve the recognition results of automatic speech recognition, speaker adaptation is used. Speaker adaptation allows to incorporate individual speaker characteristics in both speaker dependent and speaker independent references. A method and a device for continuously updating existing references is known from WO 95/09416. The method and the device described in WO 95/09416 allow to adapt existing references to changes in a speaker's voice and to changing background noise. An adaptation of an existing reference in accordance with a spoken utterance takes place each time a recognition result which corresponds to an existing reference is obtained, i. e., each time a spoken utterance is recognized.

It has been found that speaker adaptation of existing references generally improves the accuracy of automatic speech recognition. However, the accuracy of automatic speech recognition using continuously adapted references generally shows fluctuations. This means that the recognition accuracy does not continuously improve with each adaptation process. To the contrary, the recognition accuracy may also temporarily decrease.

There is, therefore, a need for a method and a device for generating an adapted reference for automatic speech

recognition which is less prone to a deterioration of the recognition accuracy.

5

SUMMARY OF THE INVENTION

The present invention satisfies this need by providing a method for generating a reference adapted to an individual speaker or a speaker's acoustic environment which comprises performing
10 automatic speech recognition based on a spoken utterance and obtaining a recognition result which corresponds to a currently valid reference, adapting the currently valid reference in accordance with the utterance, assessing the adapted reference and deciding if the adapted reference is used for further
15 recognition. A device for generating a speaker adapted reference which satisfies this need comprises means for performing recognition based on a spoken utterance and for obtaining a recognition result which corresponds to a currently valid reference, means for adapting the currently valid
20 reference in accordance with the utterance and means for assessing the adapted reference and for deciding if the adapted reference is used for further recognition.

According to the invention, an assessment step is conducted
25 after the currently valid speaker dependent or speaker independent reference is adapted. Then, and based on the result of the assessment, it is decided whether or not the adapted reference is used for further recognition. It can thus be avoided that a reference which is adapted in the wrong
30 direction will be used for further recognition. A deterioration of the recognition accuracy would take place e. g. if the recognition result does not match the spoken utterance and the adaptation process is carried out without an assessment step. Consequently, according to the invention, corrupted adapted
35 references can be rejected and e.g. permanent storing of corrupted adapted references can be aborted. This is preferably done prior to recognizing the next spoken utterance.

According to a first embodiment of the invention, the adapted reference is assessed by determining a distance between the adapted reference and the currently valid reference. The determined distance can then be used as input for the decision if the adapted reference is used for further recognition. Moreover, when assessing the adapted reference, the distances between the adapted reference and all other currently valid references which do not correspond to the recognition result may additionally be taken into account. This enables to base the decision if the adapted reference is used for further recognition on the question which of all currently valid references is closest to the adapted reference. Thus, if one of the currently valid references which does not correspond to the recognition result has a smaller distance from the adapted reference than the currently valid reference which corresponds to the recognition result, the adapted reference can be discarded.

According to a second embodiment of the invention, the assessment of the adapted reference is based on an evaluation of the user behaviour. The recognition result will most likely be wrong if e.g. a specific action automatically initiated upon obtaining the recognition result is immediately cancelled by the user or if the user refuses to confirm the recognition result. In this case it can automatically be decided that the adapted reference is to be discarded and not to be used for further automatic speech recognition. An evaluation of the user behaviour can also be carried out prior to adapting a currently valid reference in accordance with a user utterance. The adaptation process is thus not initiated if the user behaviour indicates that the recognition result is wrong. In other words, the adaptation only takes place if the user behaviour indicates that the recognition result is correct.

According to a third embodiment of the invention, the adapted reference is assessed both by evaluating the user behaviour and by determining a distance between the adapted reference and one or more currently valid references.

The decision if the adapted reference is to be used for further recognition may be a "hard" decision or a "soft" decision. The "hard" decision may be based on the question whether or not a certain threshold of a specific parameter obtained by assessing the adapted reference has been exceeded. On the other hand, a "soft" decision may be made by means of e. g. a neuronal network which takes into account a plurality of parameters like a specific user behaviour and a distance between an adapted reference and one or more currently valid references.

In case the decision is based on the distances between the adapted reference and one or more existing references, the parameters relevant for assessing an adapted reference are preferably obtained by analyzing a histogram of previously determined distances. The distances between the adapted reference and one or more existing references are preferably calculated with dynamic programming.

If it is decided that the adapted reference is used for further recognition, the adapted reference can be stored. In regard to storing the adapted reference several strategies can be applied. According to the most simple embodiment, an adapted reference is simply substituted for the corresponding previously valid reference. According to a further embodiment, a set of adapted references is created which is used in addition to a set of currently valid references. Thus, the currently valid references constitute a fallback position in case the adapted references are e. g. developed in a wrong direction. A further fallback position can be created by additionally and permanently storing a set of so-called mother references which constitute a set of initially created references.

Preferably, a currently valid reference is not adapted automatically after a recognition result is obtained but only after confirmation of the recognition result. The nature of the confirmation depends on the purpose of the automatic speech

recognition. If the automatic speech recognition is e. g. employed for addressing an entry in a telephonebook of a mobile telephone, the confirmation can be a setting up of a call based on the recognized entry.

5

The device for generating an adapted reference according to the invention can comprise at least two separate storing means for storing two separate sets of references. The provision of a plurality of storing means allows to delay the decision whether or not to store an adapted reference permanently. For each utterance recognizable by the means for performing the recognition a pair of references can be stored such that a first of the two references is stored in the first storing means and a second of the two references in the second storing means. Also, third storing means can be provided for storing mother references.

10

15

In connection with the plurality of storing means pointing means can be employed which set pointers that allow to determine all references currently valid for recognition of spoken utterances. According to a first embodiment, a pointer is set to this reference of each pair of references which is currently valid, i. e. which constitutes a reference to be used for automatic speech recognition. Upon generating a newly adapted reference, the reference of a pair of references to which the pointer is not set may be then overwritten by the newly adapted reference. According to a second embodiment, a pointer is set to a first storing means containing a currently valid set of references. Prior to or after an adapted reference is created, the content of the first storing means is copied in second storing means. Then, the adapted reference is stored in the second storing means such that a corresponding reference in the second storing means is overwritten by the adapted reference. If, after an assessment step, it is decided to use the adapted reference for further recognition, the pointer is shifted from the first storing means to the second storing means. Otherwise, if it is decided to discard the adapted reference, the pointer is not shifted.

20

25

30

35

A further aspect of the invention relates to a computer program product with program code means for performing the generation of an adapted reference for automatic speech recognition when the computer program product is executed in a computing unit. Preferably, the computer program product is stored on a computer-readable recording medium.

BRIEF DESCRIPTION OF THE DRAWINGS

Further aspects and advantages of the invention will become apparent upon reading the following detailed description of a preferred embodiment of the invention and upon reference to the drawings in which:

Fig. 1 is a schematic diagram of a device for generating a speaker adapted reference according to the invention;

Fig. 2 is a flow chart depicting a method for generating a speaker adapted reference according to the invention; and

Fig. 3 shows histograms of distances between references for correct and erroneous recognition results.

DESCRIPTION OF A PREFERRED EMBODIMENT

In Fig. 1, a schematic diagram of an embodiment of a device 100 according to the invention for generating a speaker adapted reference for automatic speech recognition is illustrated. The device 100 depicted in Fig. 1 can e. g. be a mobile telephone with a voice interface allowing to address a telephone book entry by uttering a person's proper name.

The device 100 comprises recognition means 110 for performing recognition based on a spoken utterance and for obtaining a recognition result which corresponds to an already existing and

currently valid reference. The recognition means 110 communicate with first 120 and second 130 storing means and with adaptation means 140 for adapting an existing reference in accordance with the spoken utterance. The first 120 and second 5 130 storing means constitute the vocabulary of the recognition means 110. A pair of references is stored for each item of the vocabulary such that a first reference of each pair of reference is stored in the first storing means 120 and a second reference of each pair of references is stored in the second 10 storing means 130. However, only one of the two reference, which constitute a pair of references is defined as currently valid and can be used by the recognition means 110. A first pointer (*) is set to the currently valid reference of each pair of references.

15 The adaptation means 140 communicate with the first 120 and second 130 storing means as well as with assessing means 150. The assessing means 150 assess the adapted reference and decide if the adapted reference is used for further recognition. The 20 assessing means 150 communicate with pointing means 160 which communicate with the first 120 and second 130 storing means.

The device 100 depicted in Fig. 1 further comprises third storing means 170 for storing a mother reference for each item 25 of the vocabulary of the recognition means 110. The mother references stored in the third storing means 170 are initially created speaker dependent or speaker independent references. The mother references can be used both in parallel to the references stored in the first 120 and the second 130 storing 30 means or as a fallback position.

Referring now to the flow chart of Fig. 2, the function of the device 100 is described in more detail.

35 Upon receipt of a signal corresponding to a spoken utterance, the recognition means 110 perform a recognition process. The signal can be provided e. g. from a microphone or a different signal source. The recognition process comprises matching a

pattern like one or more feature vectors of the spoken utterance with corresponding patterns of all currently valid references stored in either the first 120 or the second 130 storing means, i.e., all references to which the first pointers (*) are set. The first pointers (*) were set by the pointing means 160 as will be described below in more detail. If the pattern of the spoken utterance matches a pattern of a currently valid reference, a recognition result which corresponds to the matching reference is obtained by the recognition means 110. The recognition result is formed by a second pointer (>) which points to the pair of references which contains the currently valid reference matched by the spoken utterance.

After the recognition result is obtained, the telephone book entry, e.g. a person's proper name, which corresponds to the recognition result may be acoustically or optically output via an output unit not depicted in Fig. 1 to a user for confirmation. After the user confirms that the output is correct, e. g. by setting up the call to the person corresponding to the telephone book entry, the recognition result is output to the adaptation means 140.

The adaptation means 140 load the valid reference of the pair of references to which the second pointer (>) is set. The adaptation means 140 then adapt the loaded reference in accordance with the spoken utterance. This can e. g. be done by shifting the feature vectors of the loaded reference slightly towards the feature vectors of the spoken utterance. Thus, an adapted reference is generated based on both the loaded reference and the spoken utterance. After the adapted reference has been generated, the adaptation means 140 store the adapted reference by overwriting the non-valid reference of the pair of references to which the second pointer (>) is set.

After the adapted reference has been stored, the assessing means 150 assess the adapted reference as will be described below. The assessing is necessary since the user confirmation

might have been wrong. If the user has e.g. erroneously confirmed a wrong recognition result, the adapted reference would be corrupted. The adapted reference has therefore to be assessed. Based on the result of the assessment, the assessing means 150 decide if the adapted reference is used for further recognition.

If it is decided to use the adapted reference for further recognition, the first pointer (*) is shifted within the current pair of references to the adapted reference by the pointing means 160. Consequently, the newly adapted reference to which the first pointer (*) has been set will constitute the valid reference in terms of recognizing a subsequent spoken utterance. The adapted reference is thus stored permanently. On the other hand, if it is decided to reject the adapted reference, the position of the first pointer (*) is not changed within the pair of references to which the second pointer (>) is set. Consequently, the adapted reference may be overwritten in a subsequent recognition step.

In the following, the function of the assessing means 150 is described by means of two exemplary embodiments.

According to a first embodiment, the assessing means 150 determine the distance between the newly adapted reference and the currently valid reference of the pair of references to which the second pointer (>) is set. The distance is calculated by means of dynamic programming (dynamic time warping) as described in detail in "The Use of the One-Stage Dynamic Programming Algorithm for Connected Word Recognition", IEEE Transaction Acoustics, Speech and Signal Processing, Volume ASSP-32, No. 2, pp. 263 to 271, 1984, herewith incorporated by reference.

After the distance has been calculated, the assessing means 150 decide whether or not to use the adapted reference for further recognition, i.e., whether or not to shift the corresponding first pointer (*), based on a distance threshold. This means

that the adapted reference is only used for further recognition if the distance between the adapted reference and the corresponding currently valid reference does not exceed this threshold.

5

The distance threshold can be obtained by analyzing a histogram of previously determined distances corresponding to real life data as depicted in Fig. 3. Fig. 3 shows two different histograms 300, 310. The first histogram 300 reflects the distances between adapted references and existing references for correct recognition results and the other histogram 310 for erroneous recognition results. From Fig. 3 it becomes clear that the average distance for correct recognition results is smaller than the average distance for erroneous recognition results. In order to create the histograms depicted in Fig. 3, a large amount of data has to be analyzed.

10

The above mentioned distance threshold for deciding whether an adapted reference is used for further recognition is chosen e.g. at the crossing or in the vicinity of the crossing of the two histograms 300, 310 (as indicated by the line 320). Consequently, an adapted reference is used for further recognition if the calculated distance does not exceed the distance corresponding to the line 320.

20

25

If only the distance between the adapted reference and the corresponding currently valid reference of the pair of references to which the second pointer (>) is set is determined, only the shifting between these two references is taken into account. According to a second, exemplary embodiment for assessing the adapted reference, the distance between the adapted reference and all other currently valid references to which first pointers (*) are set are additionally calculated. Consequently, the shifting of the adapted reference with respect to all currently valid references constituting the vocabulary of the recognition means 110 is taken into account. The distances can again be calculated by means of dynamic programming.

30

35

After the distances between the adapted reference and all other currently valid references have been calculated, the currently valid reference having the smallest distance from the adapted reference is determined. Should the currently valid reference having the smallest distance from the adapted reference not correspond to the pair of references to which the second pointer (>) is set, it is decided that the adapted reference is not used for further recognition. Otherwise, should the currently valid reference with the smallest distance from the adapted reference belong to the pair of references to which the second pointer (>) is set, the adapted reference is used or further recognition.

EPO - Munich
69
07. Nov. 2000

CLAIMS

1. A method for generating an adapted reference for automatic speech recognition, comprising:

5

- performing recognition based on a spoken utterance and obtaining a recognition result which corresponds to a currently valid reference;

10

- adapting the currently valid reference in accordance with the spoken utterance; and

- assessing the adapted reference and deciding if the adapted reference is used for further recognition.

15

2. The method according to claim 1, wherein the adapted reference is assessed by determining a distance between the adapted reference and the currently valid reference.

20

3. The method according to claim 2, wherein the adapted reference is assessed by further determining distances between the adapted reference and currently valid references which do not correspond to the recognition result.

25

4. The method according to one of claims 2 and 3, further comprising analyzing a histogram of previously determined distances in order to obtain one or more parameters for deciding if the adapted reference is used for further recognition.

30

5. The method according to one of claims 2 to 4, wherein the distance or the distances are determined by dynamic programming.

35

6. The method according to one of claims 1 to 5, wherein the adapted reference is assessed based on a user behaviour.

7. The method according to one of claims 1 to 6, further comprising substituting the currently valid reference by the adapted reference.

5 8. The method according to one of claims 1 to 6, further comprising storing the adapted reference in addition to the currently valid reference.

10 9. The method according to one of claims 1 to 8, wherein the adapted reference is created only when a user behaviour indicates that the recognition result is correct.

15 10. A computer program product with program code means for performing the steps according to one of claims 1 to 9 when the product is executed in a computing unit.

20 11. The computer program product with program code means according to claim 10 stored on a computer-readable recording medium.

25 12. A device (100) for generating an adapted reference for automatic speech recognition, comprising:

25 - means (110) for performing recognition based on a spoken utterance and for obtaining a recognition result which correspond to a currently valid reference;

30 - means (140) for adapting the currently valid reference in accordance with the spoken utterance; and

35 - means (150) for assessing the adapted reference and for deciding if the adapted reference is used for further recognition.

13. The device according to claim 12, further comprising first and second storing means (120, 130) for storing a first and a second set of references.
- 5 14. The device according to claim 12 or 13, further comprising third storing means (170) for storing a set of mother references.
- 10 15. The device according to claim 13 or 14, further comprising pointing means (160) for setting pointers (*, >) which allow to determine all references currently valid for recognition of spoken utterances
- 15 16. The device according to claim 15, wherein the pointing means set a pointer to either the first or the second storing means depending on whether the first or the second set of references are to be used for recognition.

THIS PAGE BLANK (USPTO)

ABSTRACT

The invention relates to a method for generating an adapted reference for automatic speech recognition. In a first step, recognition is performed based on a spoken utterance and a recognition result which corresponds to a currently valid reference is obtained. In a second step, the currently valid reference is adapted in accordance with the utterance in order to create an adapted reference. In a third step, the adapted reference is assessed and it is decided if the adapted reference is used for further recognition.

(Fig. 2)

6022

THIS PAGE BLANK (USPTO)




Fig. 1

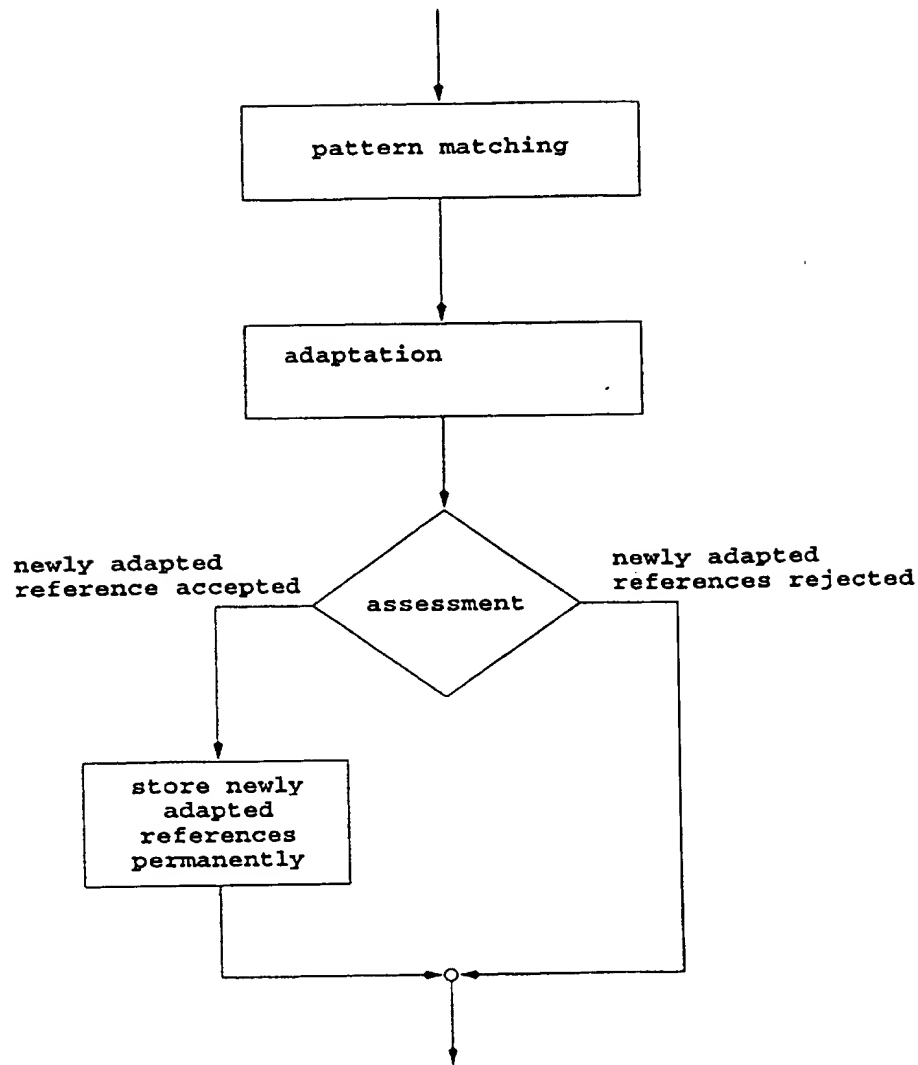


Fig. 2




Fig. 3

THIS PAGE BLANK (USPTO)